Fermilab Site Report
Fall 2013 HEPiX

Keith Chadwick
Scientific Technical Architecture Group
Fermilab

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

Fermilab

Acknowledgements

I am reporting on the accomplishments of the many individuals and organizations of the Fermilab Computing Sector,

Please credit them with all the accomplishments and assign to me any misunderstandings.

Fermilab

Some quick facts about Fermilab:

- We are located in Batavia, IL (~45 miles/~60 km west of Chicago).
- Is the premier US based HEP lab,
- Provides services to the US HEP program,
- Hosts experiments in the Energy, Intensity and Cosmic frontiers,
- Is a strong collaborator with the Open Science Grid (OSG) and the World-wide LHC Computing Grid (WLCG),
- Hosts the US Tier 1 facility for the CMS collaboration,
- And one of the homes of Scientific Linux.

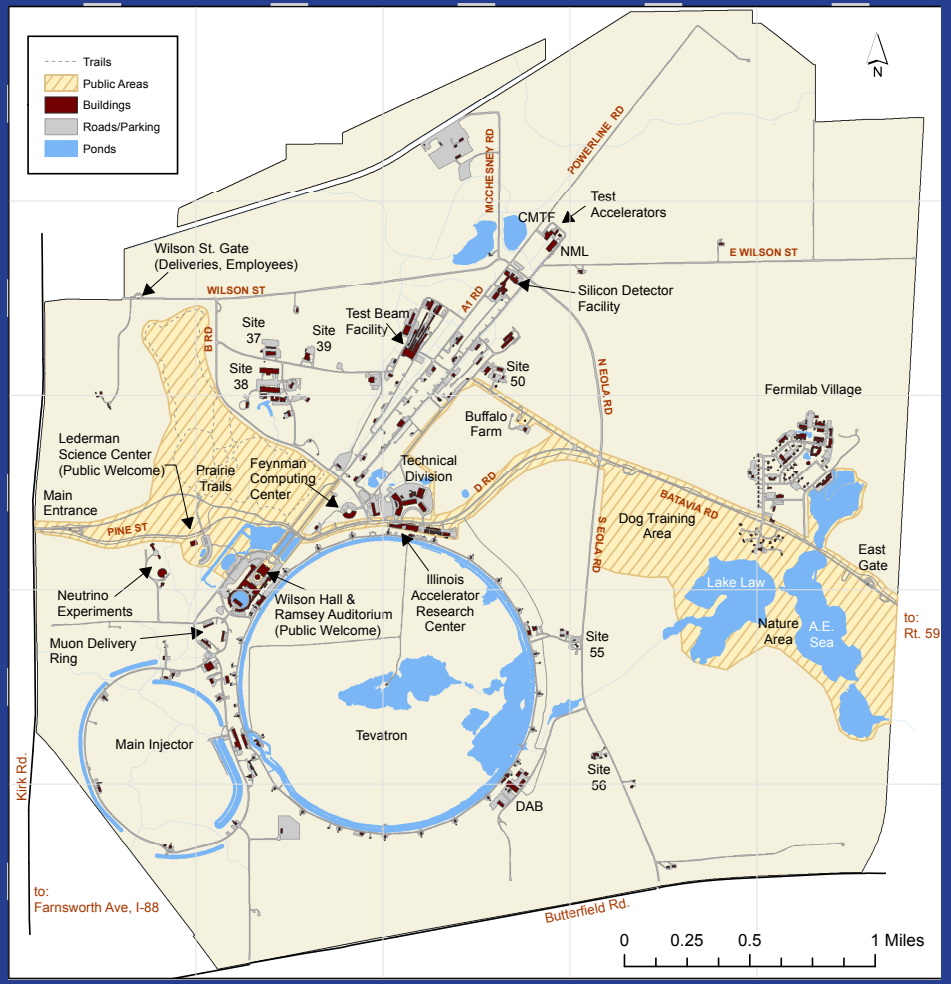

Scientific Linux

Recently mentioned in the online version of the Chicago Tribune on 28-Oct-2013.

Blue Sky Innovation
Chicago's grand contribution to science

"If the Linux operating system can be thought of as a common "language," at least one powerful "dialect" - Scientific Linux - originated at Fermilab near Chicago. Scientific Linux, No. 10 on our list of top Chicago innovations, is used at top universities, at national laboratories and on the orbiting International Space Station."

#10: FermiLinux (1998)

Since programmer Linus Torvalds created Linux as a free source code in 1991, it has spread around the world as a useful, flexible operating system that serves as an alternative to Microsoft's Windows and others. But if Linux can be thought of as a common "language," at least one powerful "dialect" originated at Fermilab near Chicago. In collaboration with the European Organization for Nuclear Research, or CERN, and other labs and universities, Fermilab recompiled a version of Linux — now called "Scientific Linux" because of its use in the global scientific community.

In 1997, Fermilab Lead Scientific Linux Developer Connie Sieh tinkered with a version of Linux distributed by RedHat Inc. — a language attractive because it was free, could be installed in batches and was simple enough for scientists to install on their own, according to Fermilab's website.

In 1998, Fermilab renamed it "FermiLinux," and it became the language used to program physics experiments. By 2004, developers from Fermilab, CERN and other collaborators had polished it up into "Scientific Linux" for use in every high-energy physics lab.

The international language of science used to be Latin.
Now it is "Scientific Linux."

More than 140,000 users run Scientific Linux, making it easier for scientists to do work while visiting other institutions. The language is "spoken" at top universities, at national laboratories, at CERN, and on the orbiting International Space Station.

When scientists experiment with the world's highest energy physics, they program the experiments in Scientific Linux. This year, one of those labs — the Large Hadron Collider at CERN — was where scientists from around the world collaborated to discover the Higgs boson, also referred to as the "God particle."

Reorganization / New Roles

The Computing Sector, Core Computing Division and Scientific Computing Division all reorganized effective 1-Jul-2013.

- Details of this reorganization are in the "extra sides" attached to this presentation.

In addition…

Vicky White is now Fermilab COO,

- She also retains her roles as CIO and Associate Director for Computing.

New Deputy CIO - Jin Chang (formerly CIO of NYC Department of the Environment).

New Groups in the Scientific Computing Division

Scientific Technical Architecture (STA):

- This group documents, captures, fosters, and promotes a coherent architecture, including standards and best practices, adopted by scientific software and computing, in conjunction with the Office of Project Management Enterprise Architecture activity.

  ⇒ New strategy on virtualization and on-demand services (see extra slides).

Entrepreneurial Ventures

(aka External Opportunities and Marketing):

- This group works with the Scientific Computing Division staff and Experiment computing organizations to identify, promote and facilitate opportunities for external scientific computing activities and collaborations and generate technical materials to market, communicate and foster entrepreneurial ventures.

ITIL Service Management

Prepping for SLA, Availability, Capacity, and annual reviews for certified services;

Starting to fully onboard scientific services:

- Data Movement and Storage is "blazing the trail";
- Previously selected scientific services were onboarded to only incident and change management.

Updating experiment TSWs [MOUs] utilizing new service oriented framework.

ITIL Roles and Personnel:

Fermilab IT Service Management

ITIL Service Management

Incident Management (snapshot from last week)

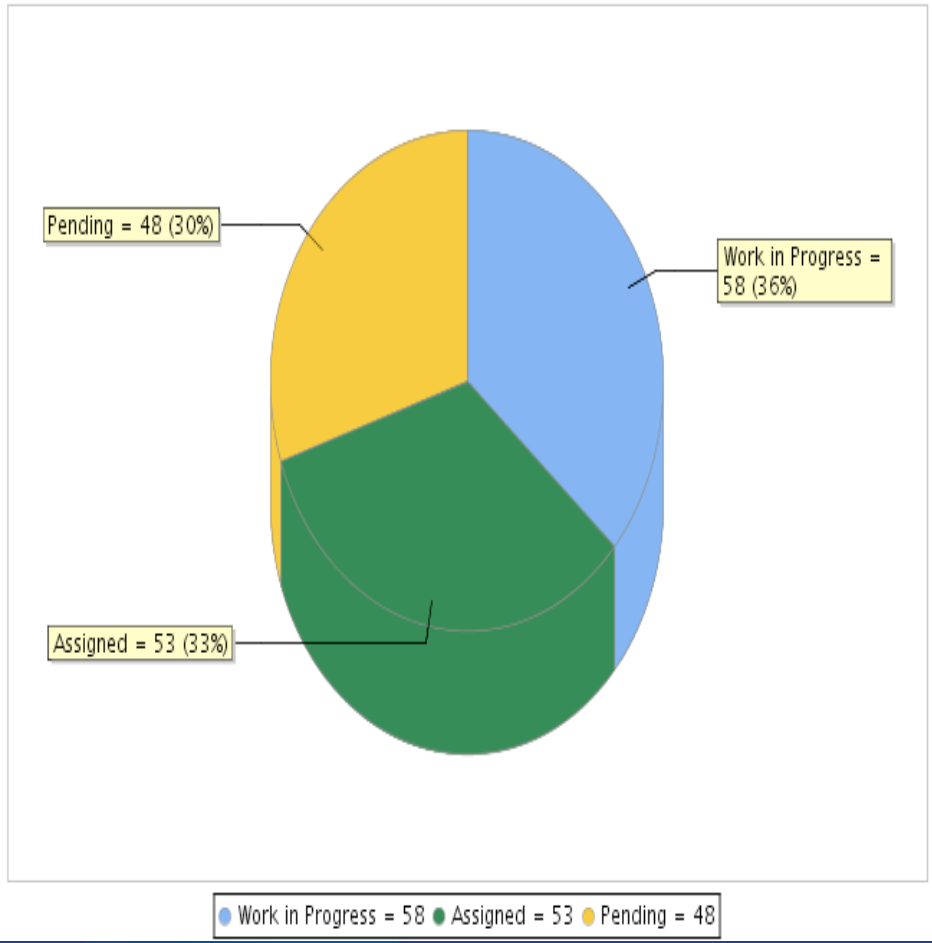

Change Management (past year)

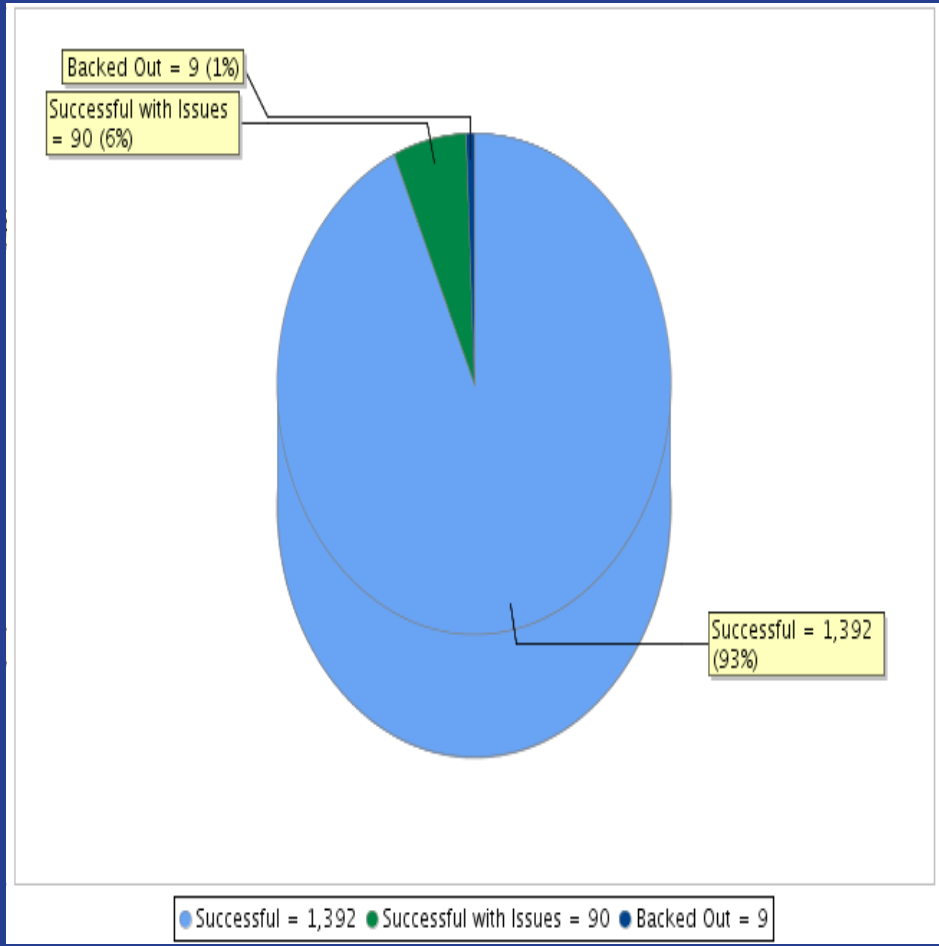

KPIs

Monthly Incident Volume

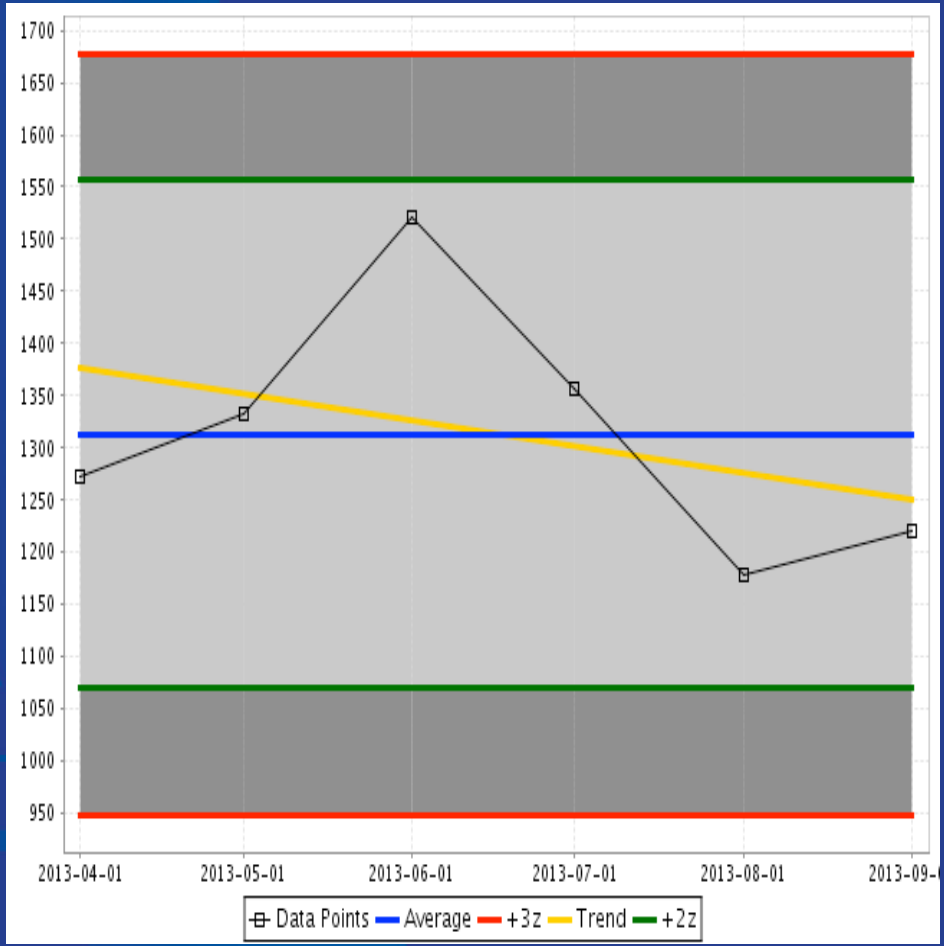

Monthly Incident Response

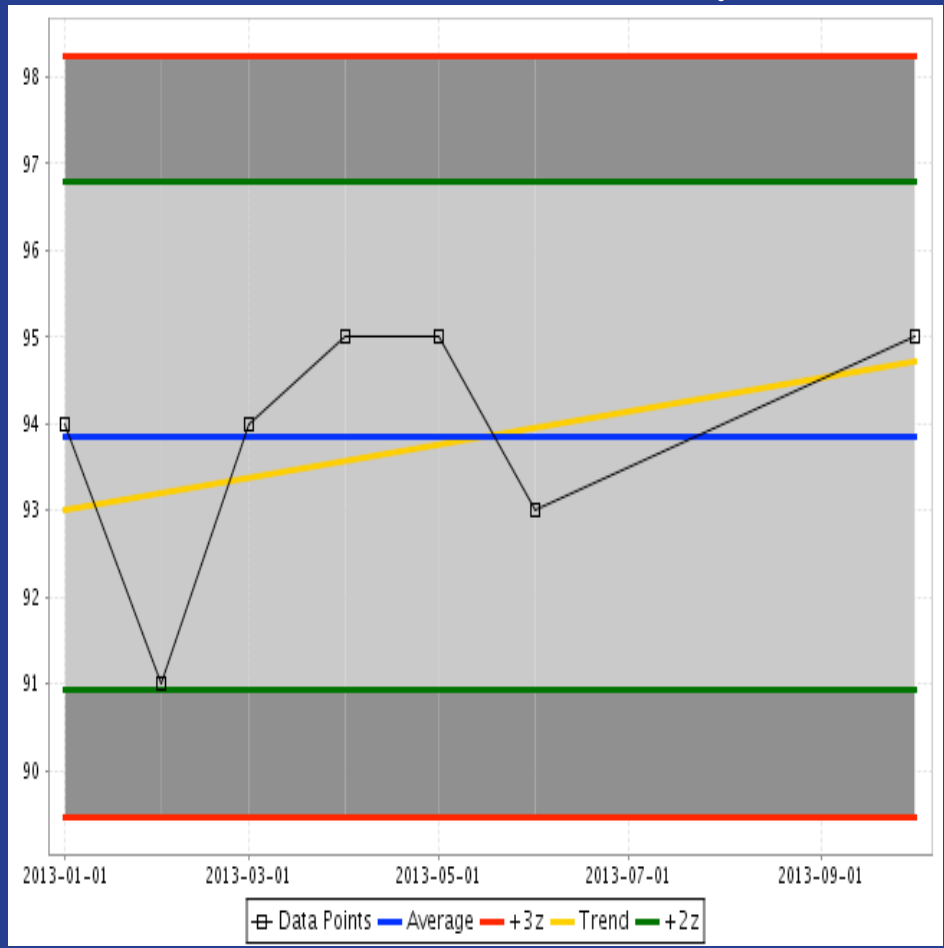

Computing Facilities

FCC:

- The FCC3 computer rooms continue to operate without any disruptions since commissioning in December 2010,

GCC:

- GCC-[A,B,C] had 8 hours of scheduled downtime,
- GCC-B and GCC-C were impacted by one "Load Shed" event in the summer of 2013,
- GCC-C also had an additional 8 hour downtime to deal with an internal electrical panel issue.

LCC:

- LCC had >99.9% uptime.

The GCC computer center received an energy star certificate for 2013.

- This is the third year in a row that this building has received the award.

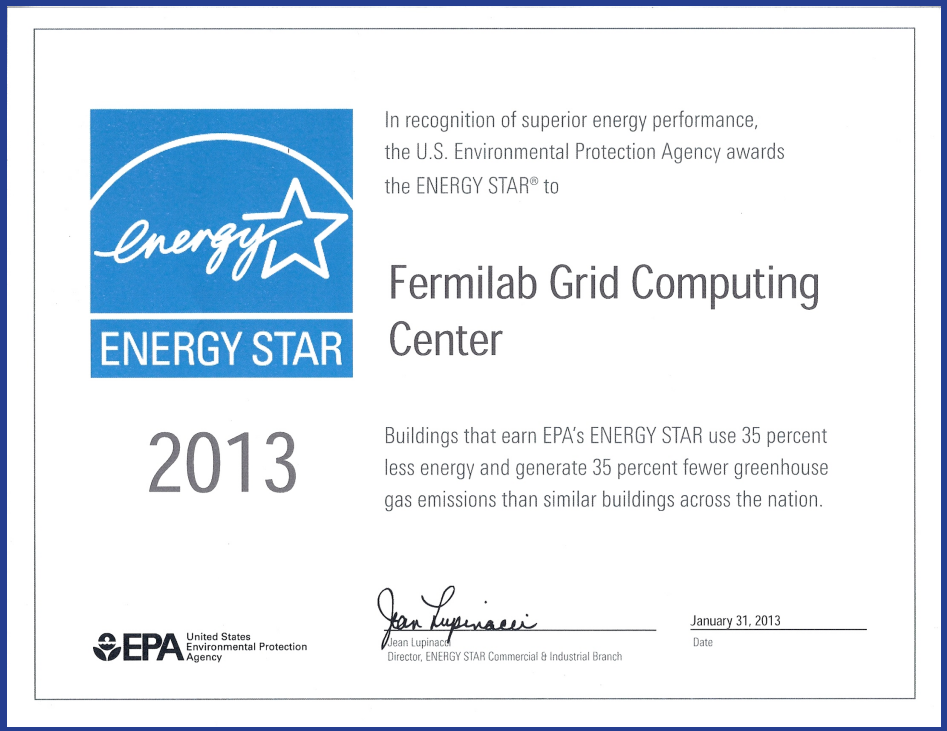

Computer Center Availability

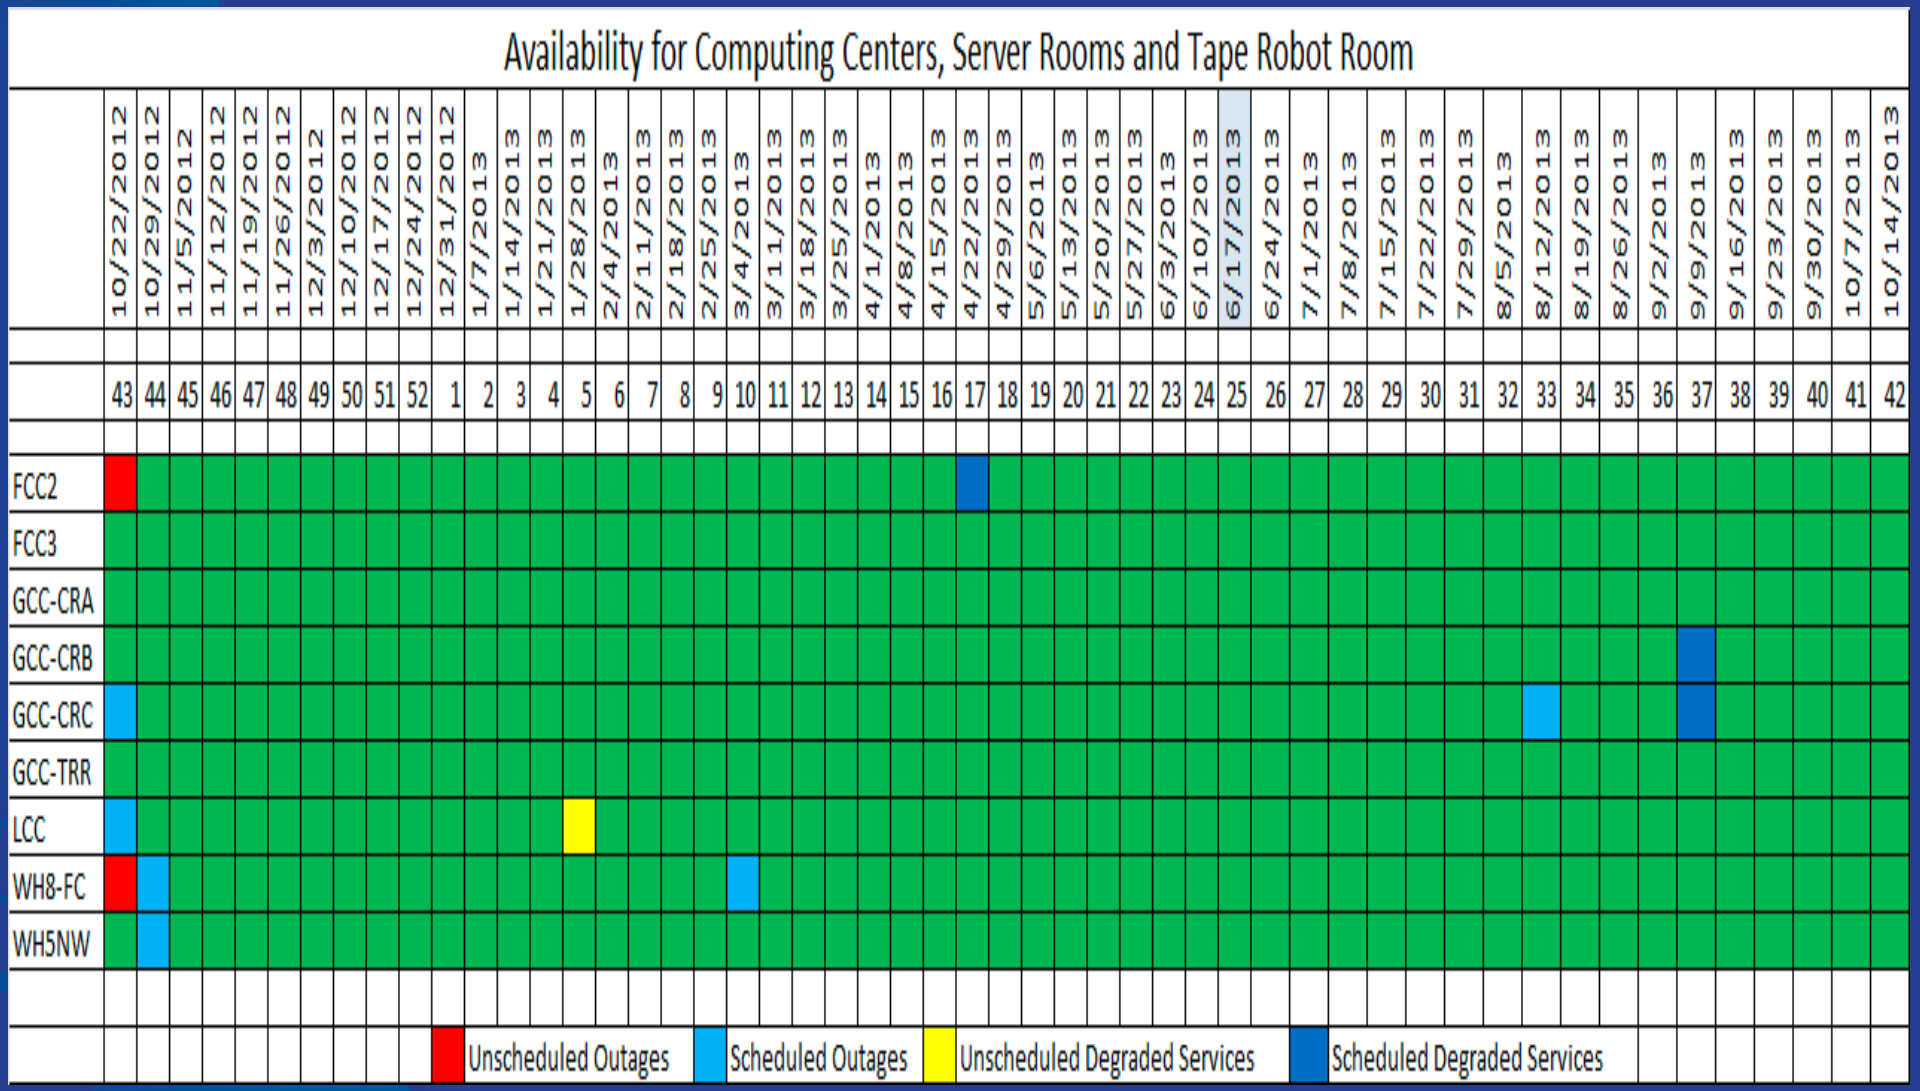

Computing (Power) Capacity

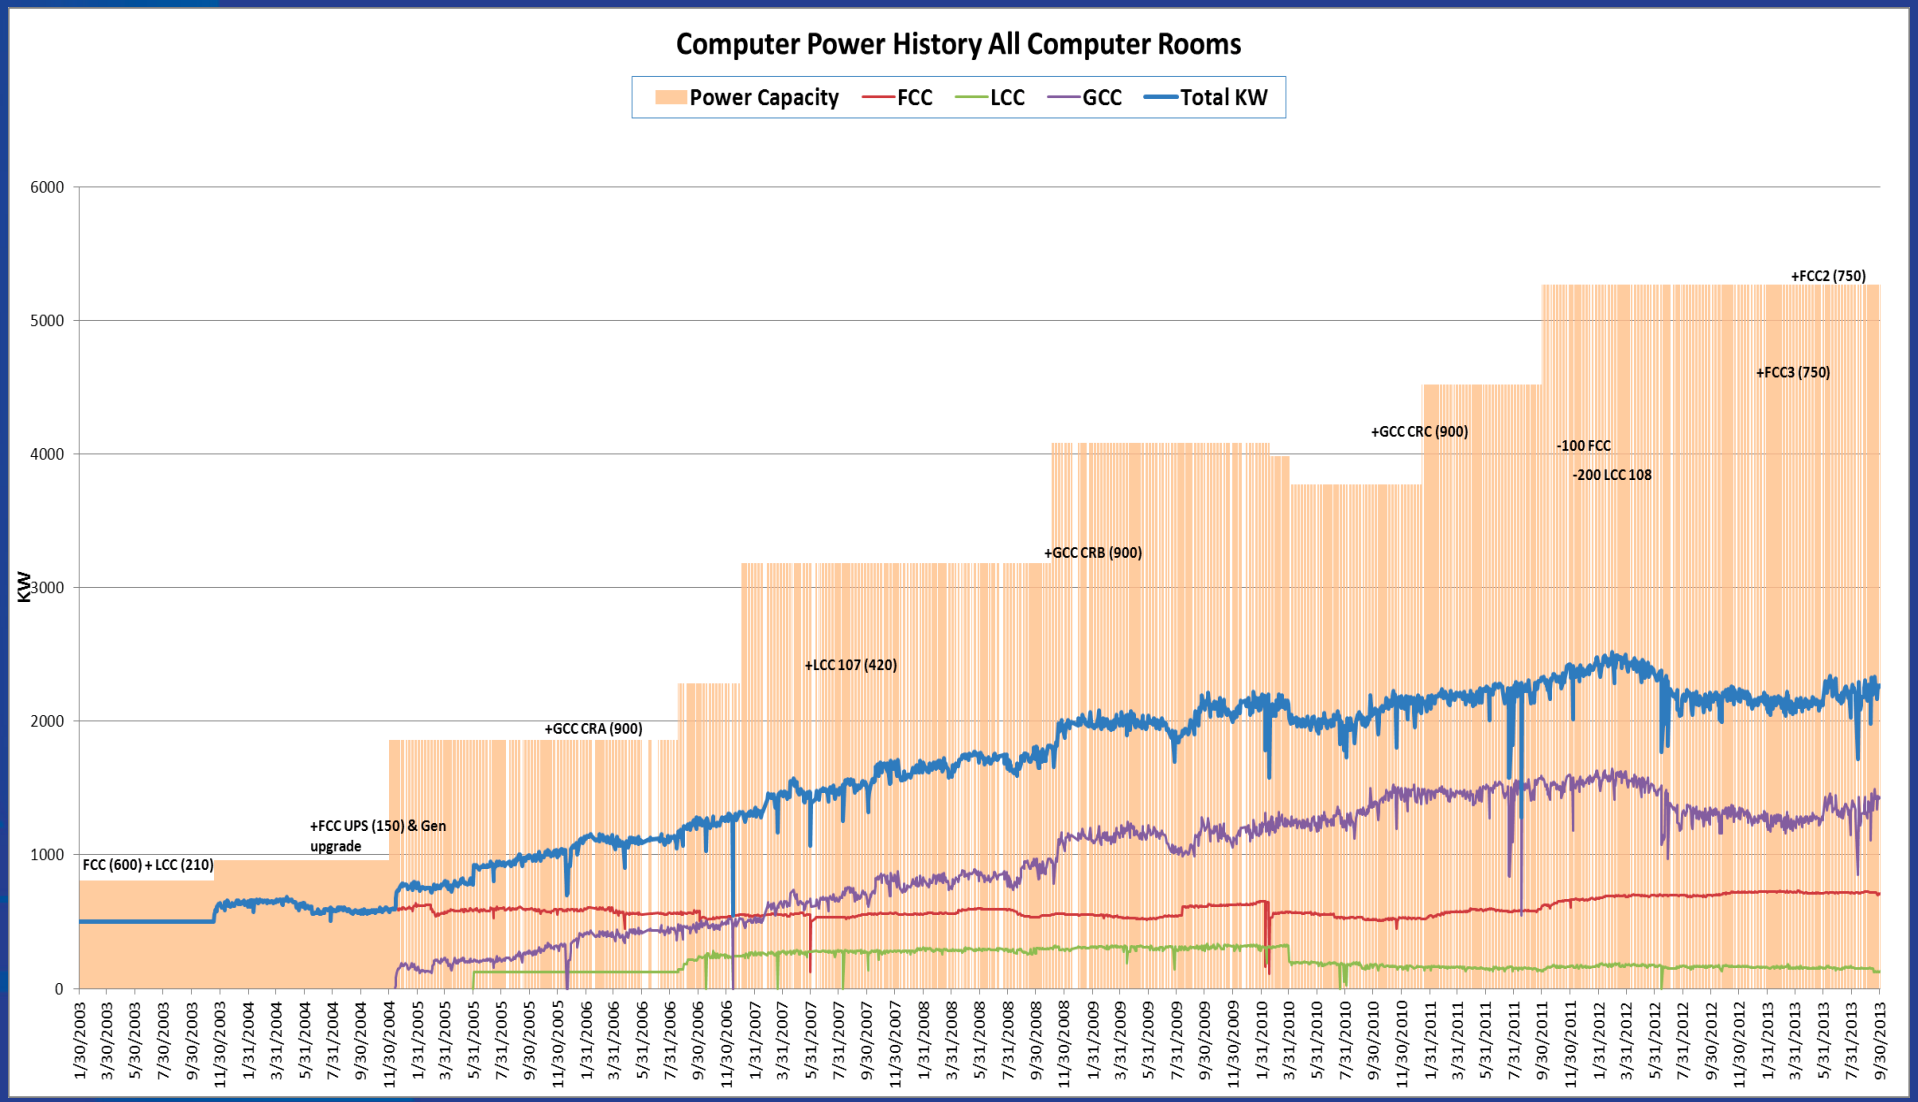

Network Initiatives

Legacy 10GE-based MAN being phased out:
- Eight 10GE channels,
- Six used for science dat.a

New 100GE-based MAN being phased in:
- One 100GE channel and three 10GE channels,
- 100GE channel is for science data movement,
- 10GE channels for general IP traffic and redundancy for science data,
- Link also being used for 100 Gb/s WAN tests (DOE 100 Gb/s testbed and transfers between Fermilab & UFL).

IPv6:
- Participation in the HEPiX IPv6 testbed uncovered a performance/configuration issue in our site infrastructure.
- Fortunately this could be resolved by a straightforward reconfiguration.

802.11ac:
- In the process of rolling out an 802.11ac wireless upgrade across the Fermilab site.

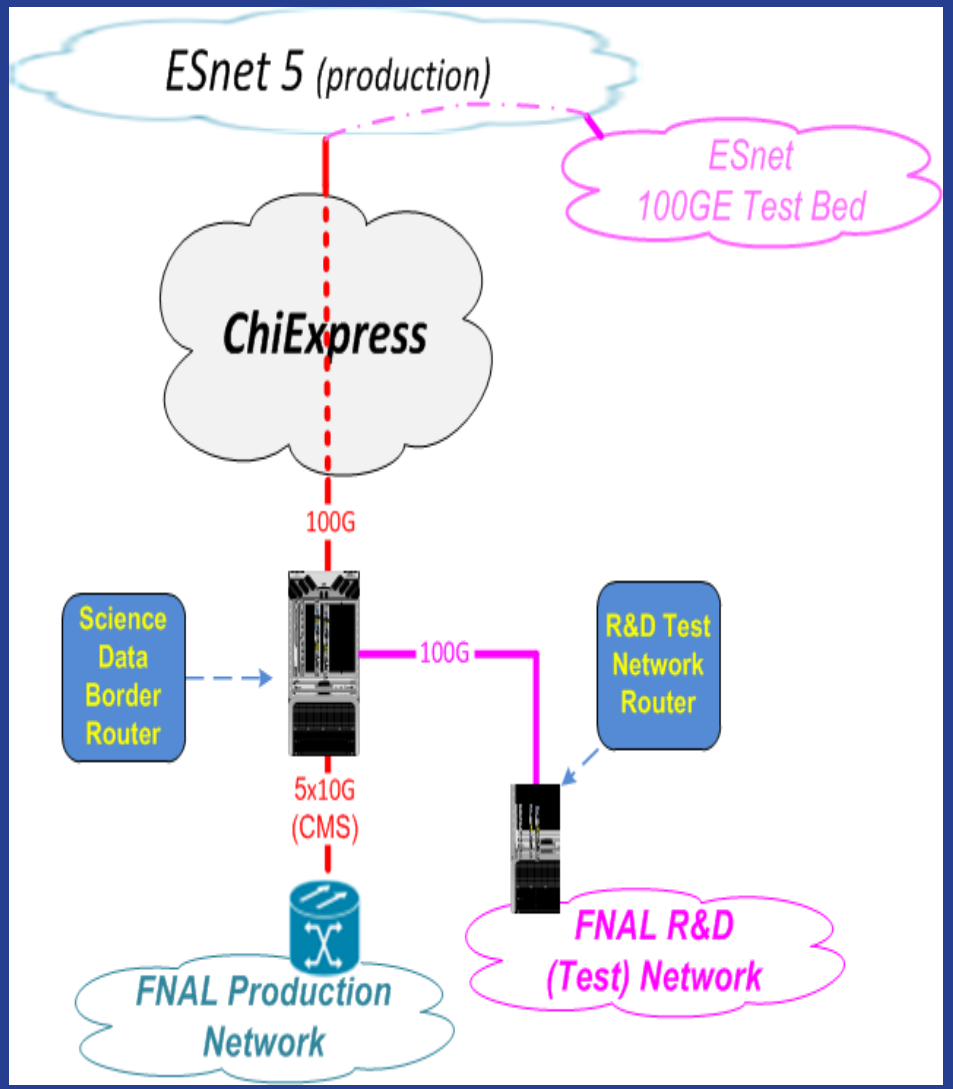

Dell Managed Services

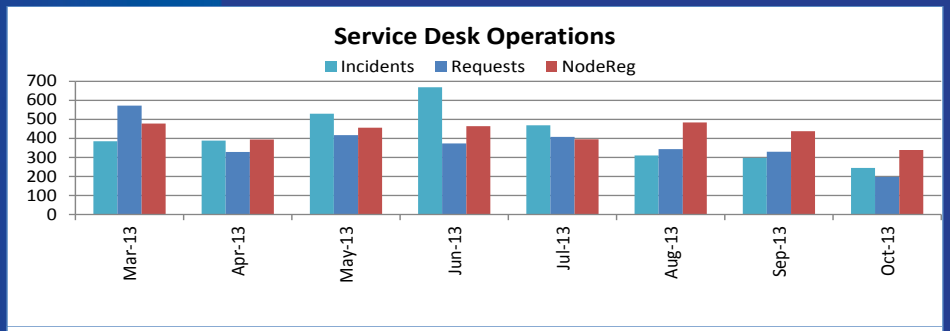

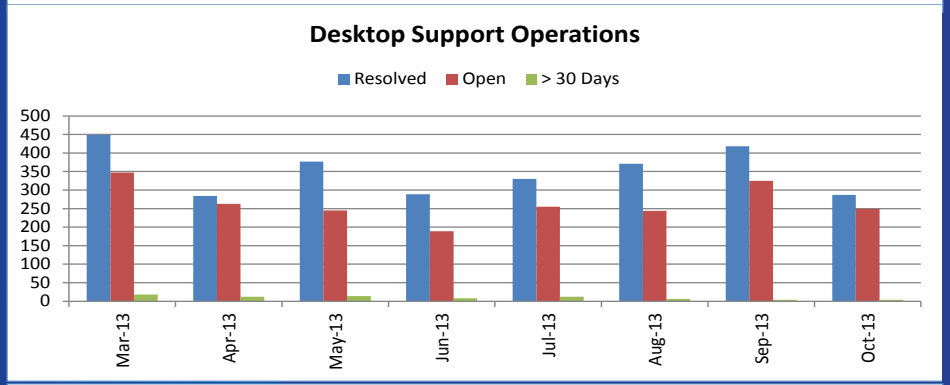

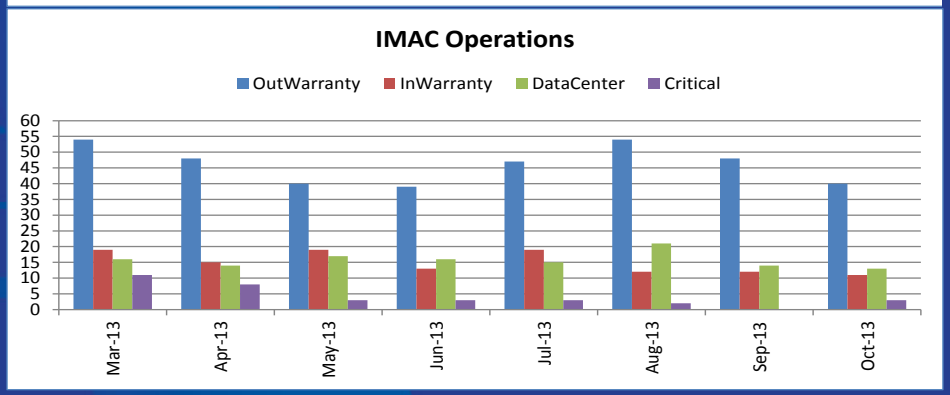

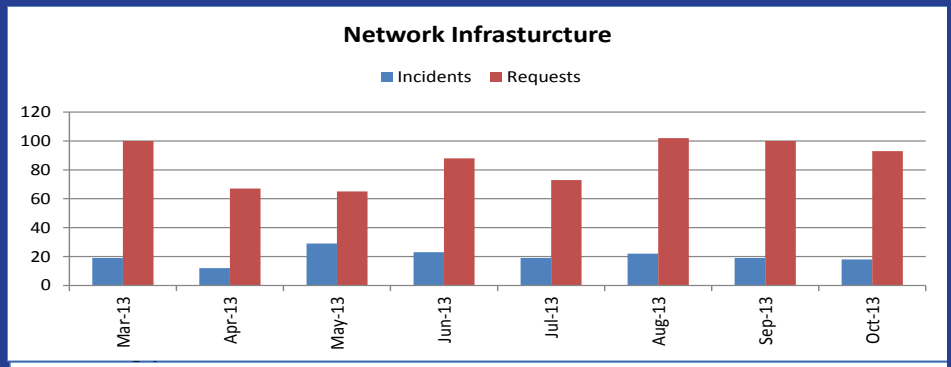

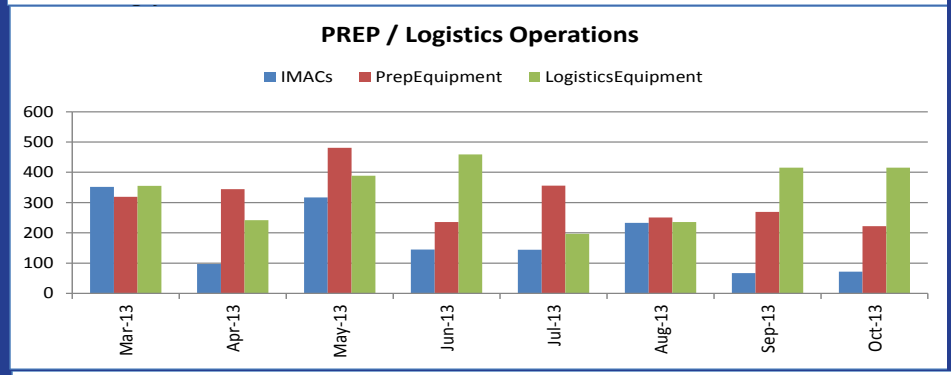

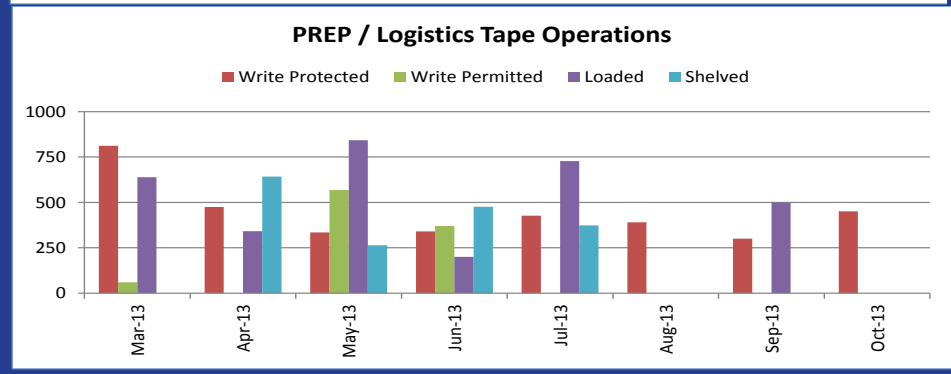

Tape Storage - Summary

- Currently have 78 PB on tape (47 PB "active")
  - Writing both LTO4 and T2 media
- Actively migrating from LTO4 to T2
  - Completed CMS, Tevatron Run II RAW,
  - 40% of Run II non-RAW completed.
- Had some T2 media contamination issues.
  - Worked closely with vendors to resolve this satisfactorily.
  - No data lost.
- Moving to denser media has left us with excess library capacity.
  - Investigating non-Fermilab archival customers.
- Increased small file aggregation capacity and added users
  - see talk by A. Moibenko later this week.

PetaBytes on Tape per FY

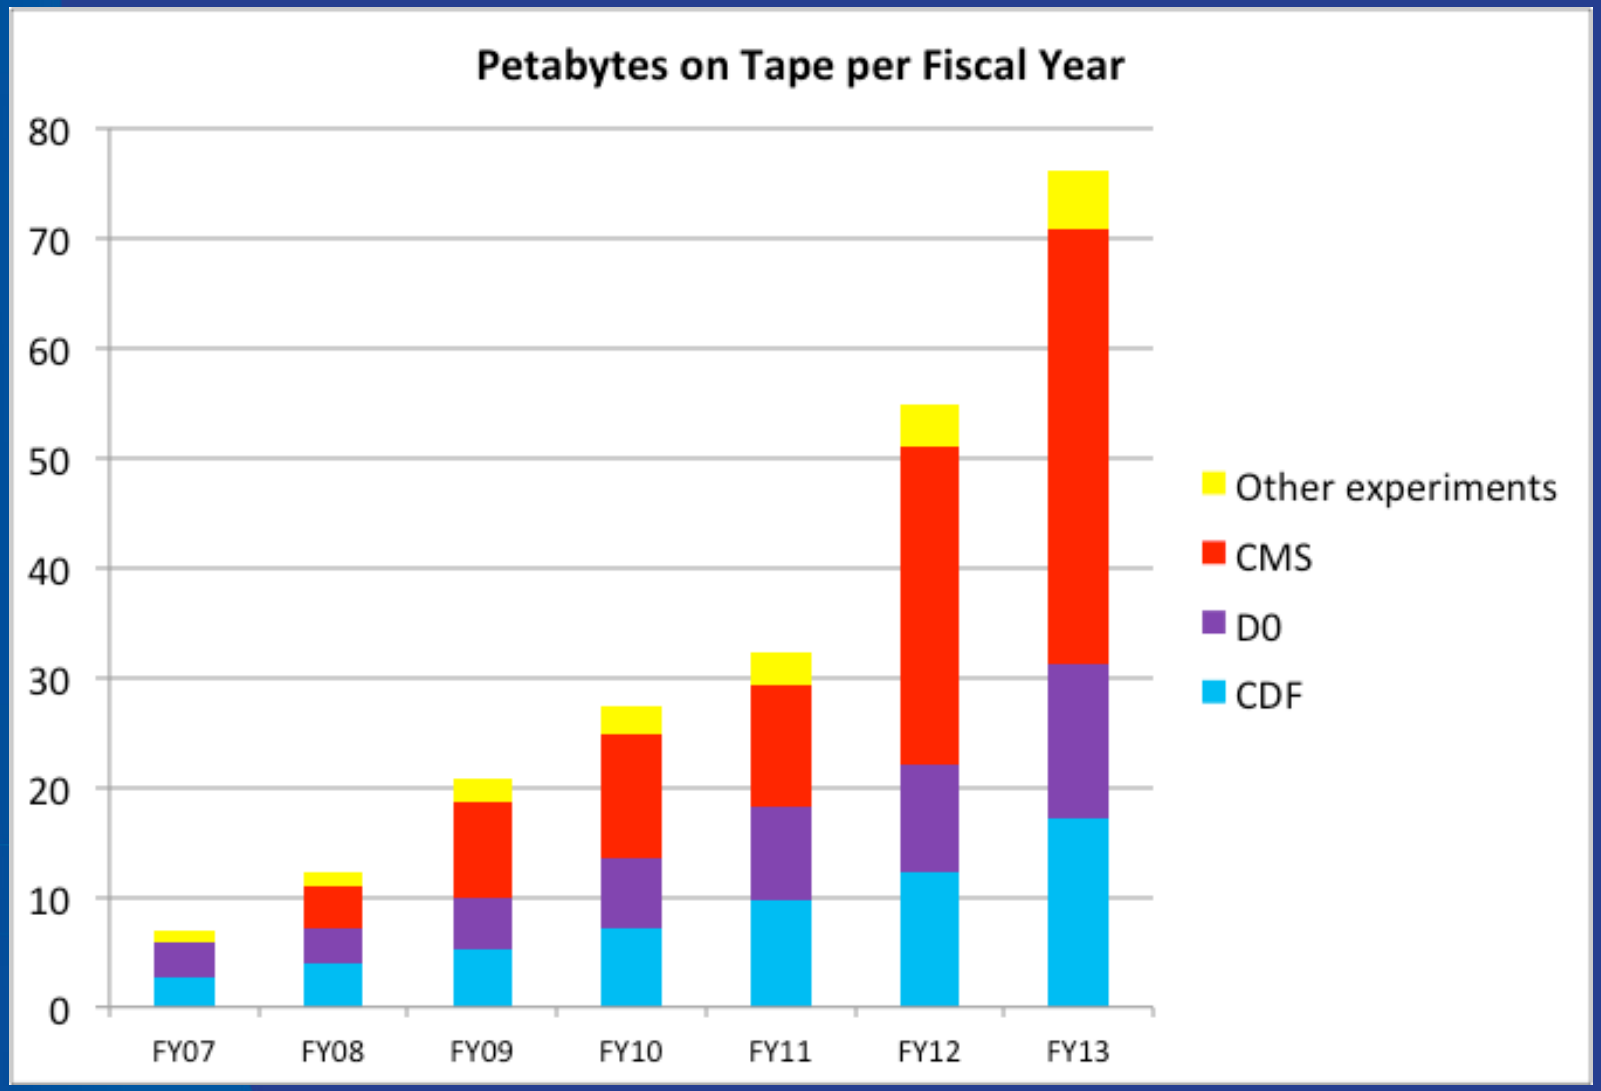

Data Movement and Storage

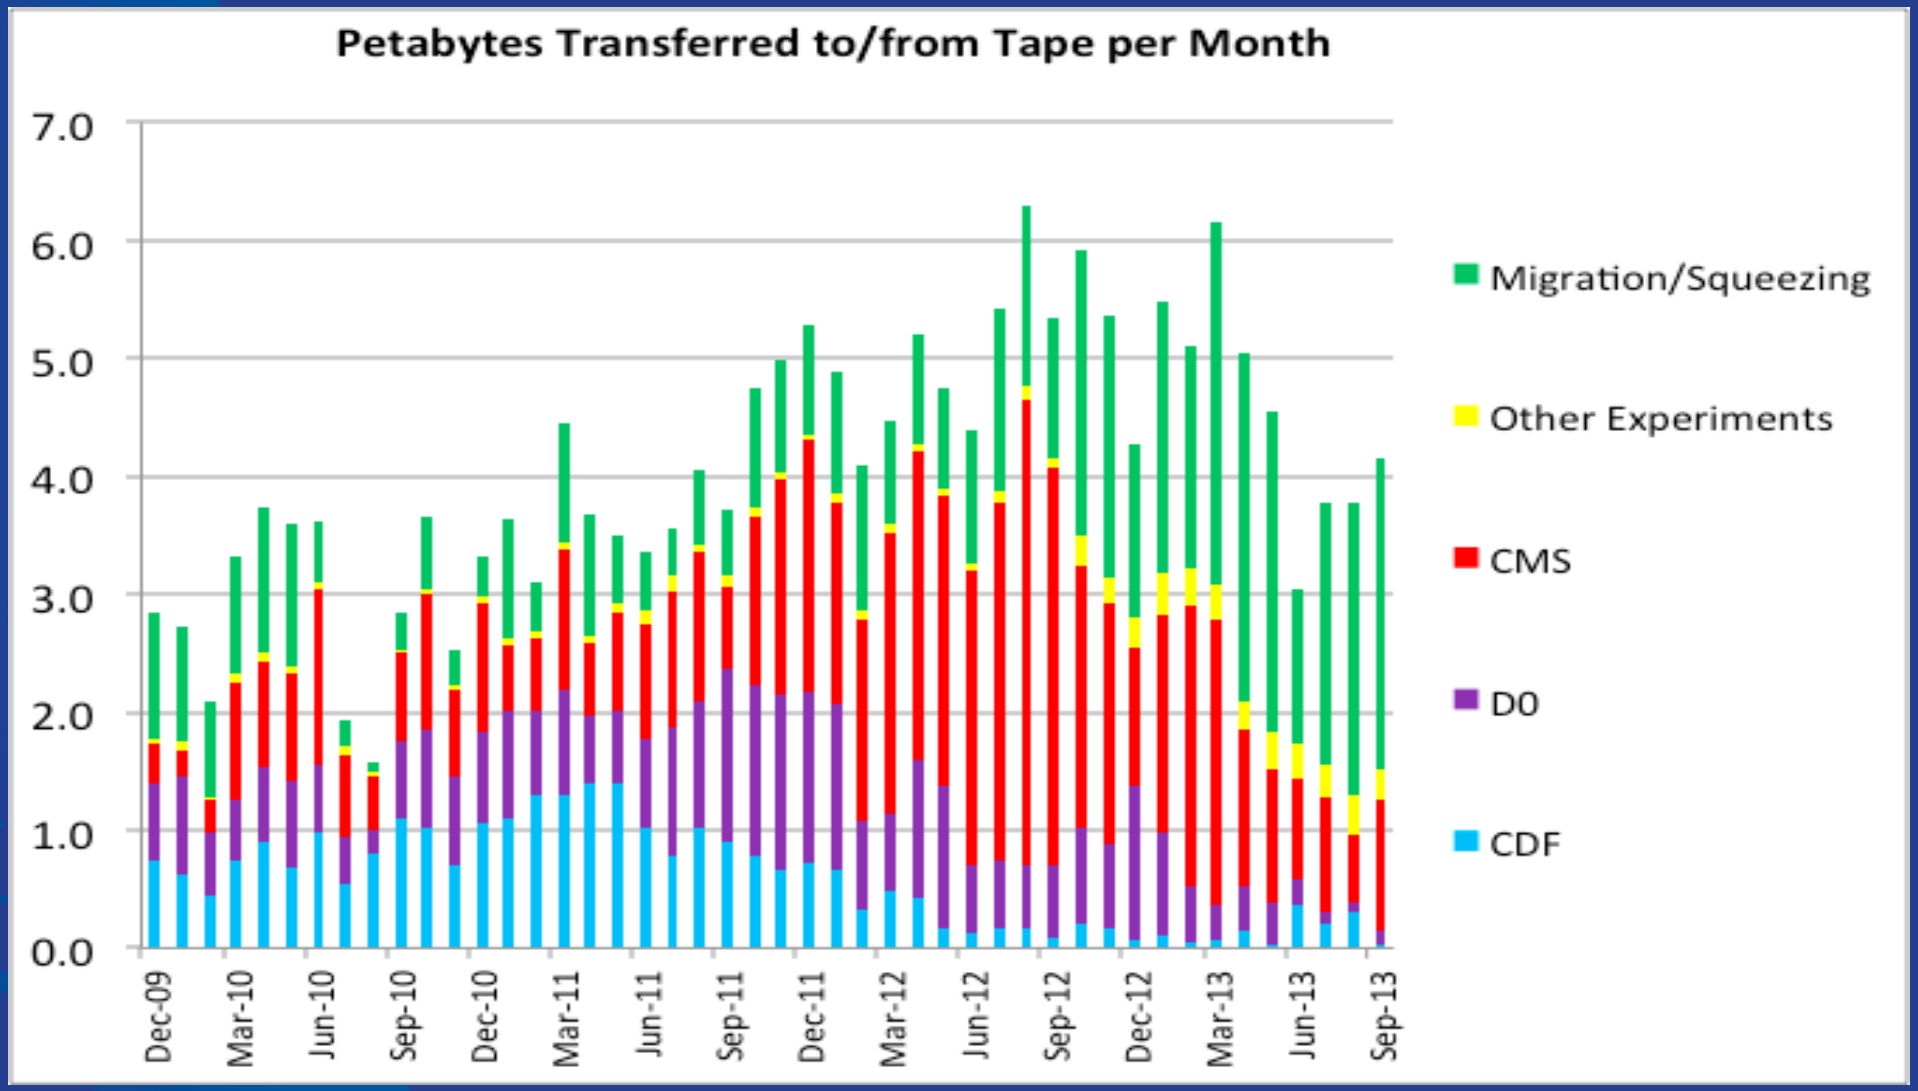

Storage Utilization by Experiment

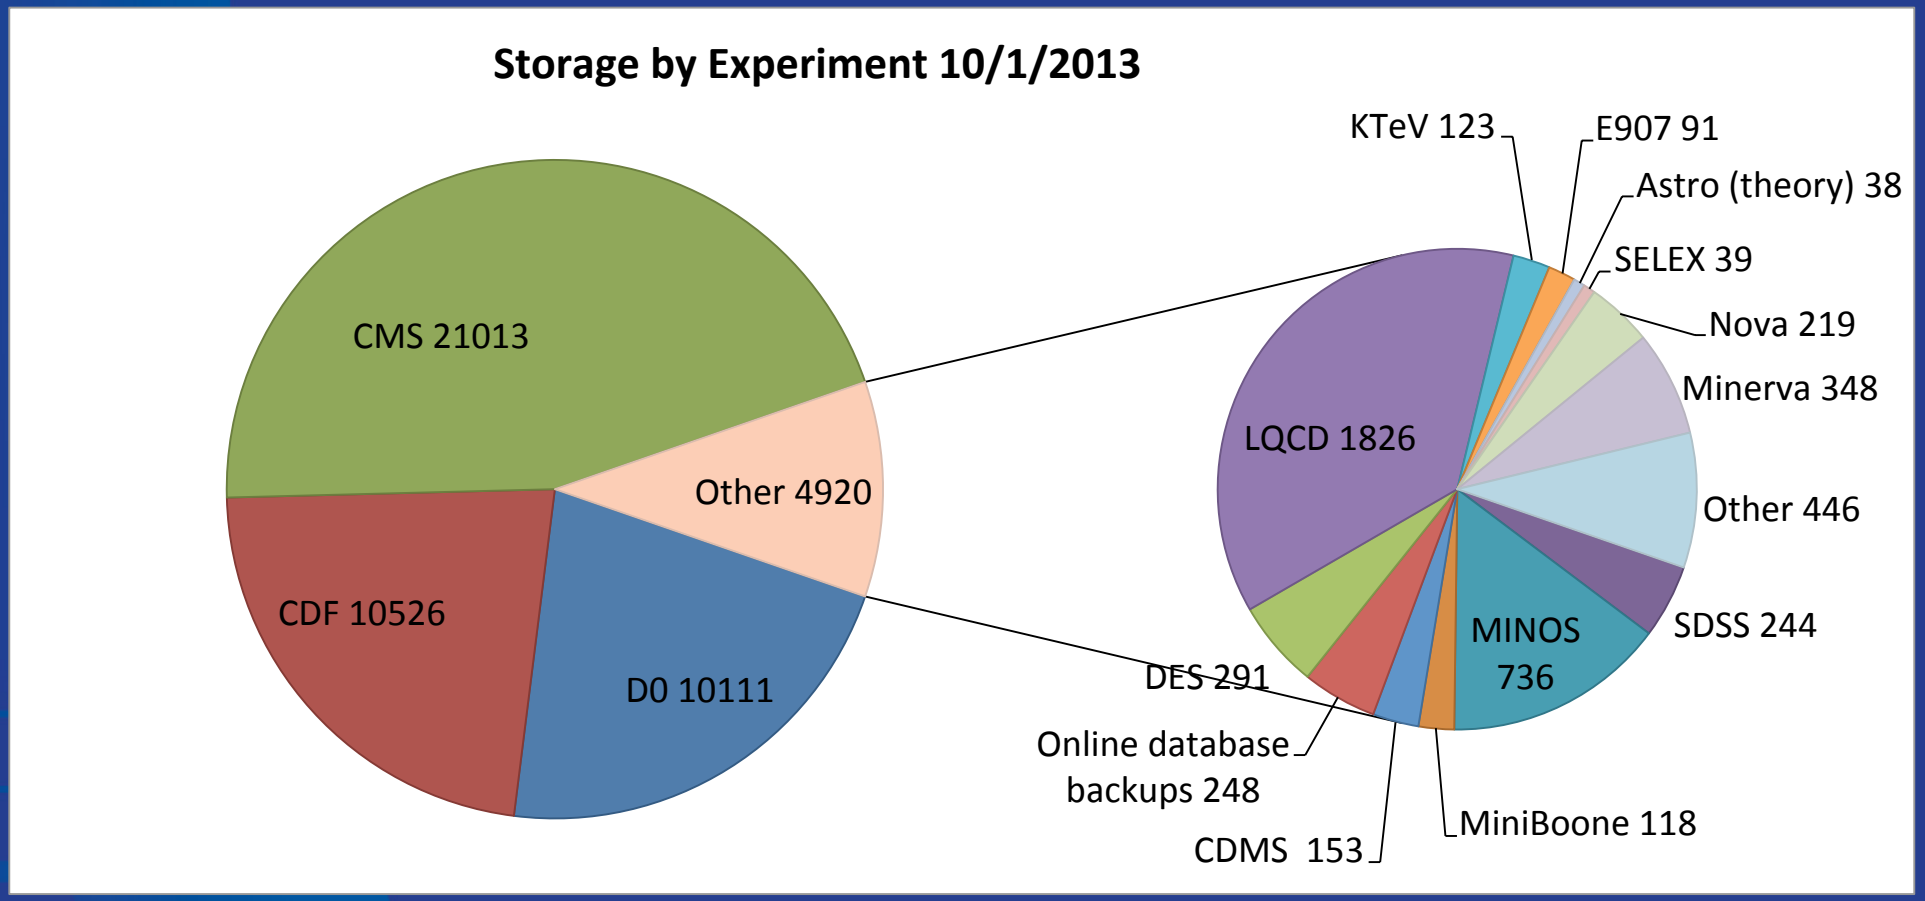

Disk Cache Storage

Deploying FY13 Purchase of ~6.5 PB:

- Nexsan E60 with 60 x 4 TB drives (raw),
- CMS deployed ~2 PB to to EOS,
- "Public" dCache (used by Intensity Frontier) upgraded from 0.18 PB to ~4 PB
  - Intensity Frontier experiments moving data access from BlueArc to dCache

CMS splitting their ~15 PB of existing dCache:

- ~13 PB non-tape backed,
- ~2 PB tape backed.

Run II:

- CDF dCache at 1.6 PB,
- D0 SAM cache at 1 PB,
- As part of their long term preservation strategy, D0 is looking into moving to dCache.

Altogether a total of approximately 24 PB disk cache.

HPC and LQCD [1]

The LQCD, Cosmology and Wilson HPC clusters are performing well.

Planning is underway for a large LQCD HPC procurement in FY2014.

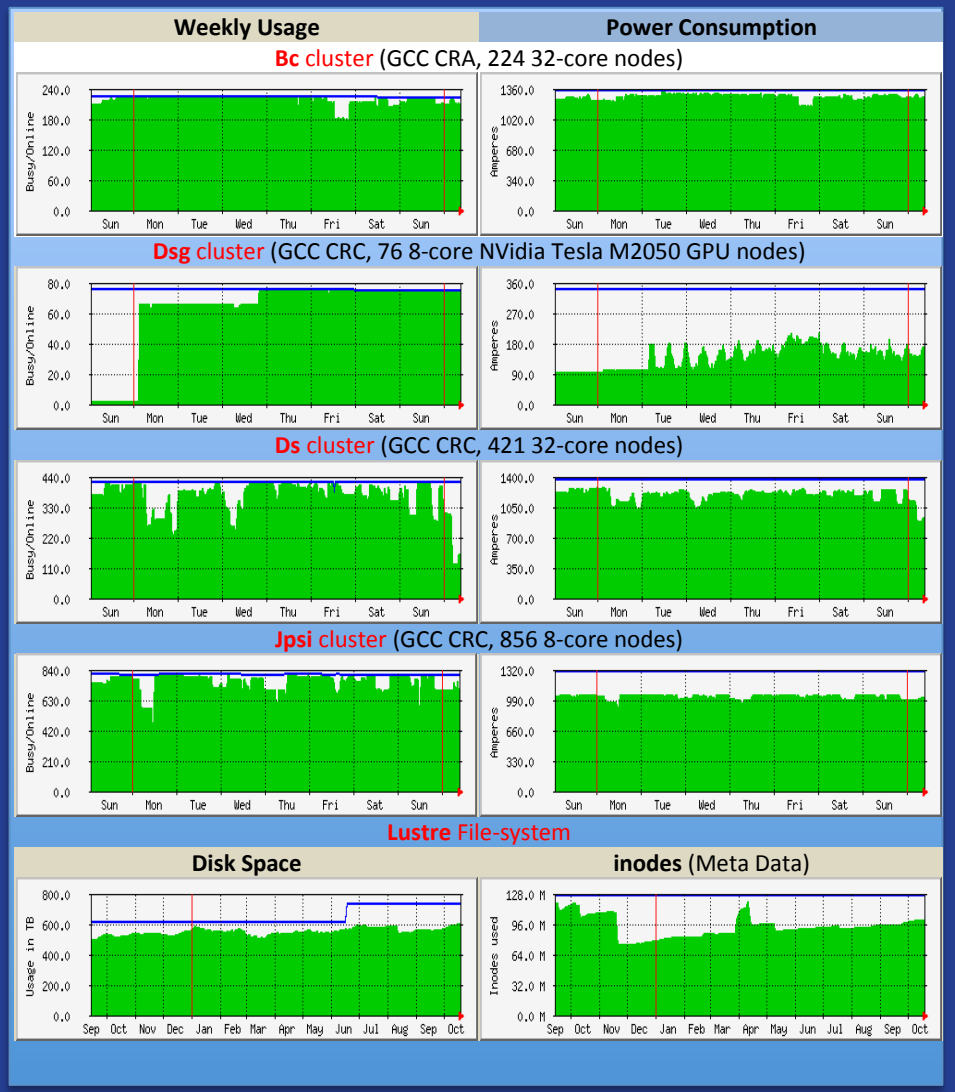

HPC and LQCD [2]

People are learning how to make effective use of new archtectures:

- an Xeon Phi based cluster,
- a NVIDIA Kepler based cluster.

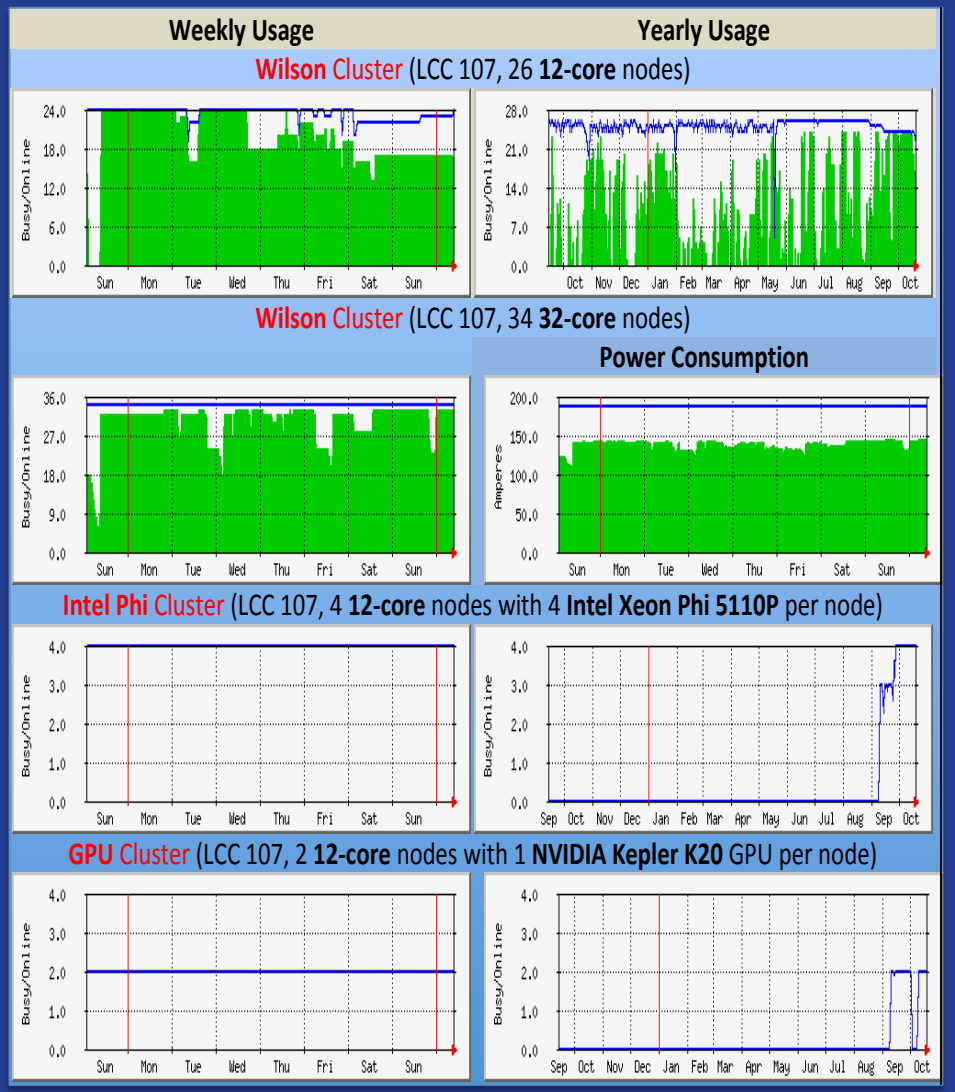

Current Fermilab Campus Grid Statistics (as of Oct 2013)

Cluster(s)	Batch System	Job Slots	Raw Occupancy	Effective Utilization
CDF	Condor	5,128	86.1	67.4
CMS T1	Condor	6,088 (*)	93.8	77.9
D0 (Merged)	PBS	5,792	74.3	54.1
GP Grid	Condor	7,522 (*)	65.1	53.0
_____		_____	_____	_____
Overall-Today		24,530	79.4	62.8
Spring HEPiX		22,934	80.4	68.3

(*) – New worker nodes being commissioned

Fermilab Campus Grid

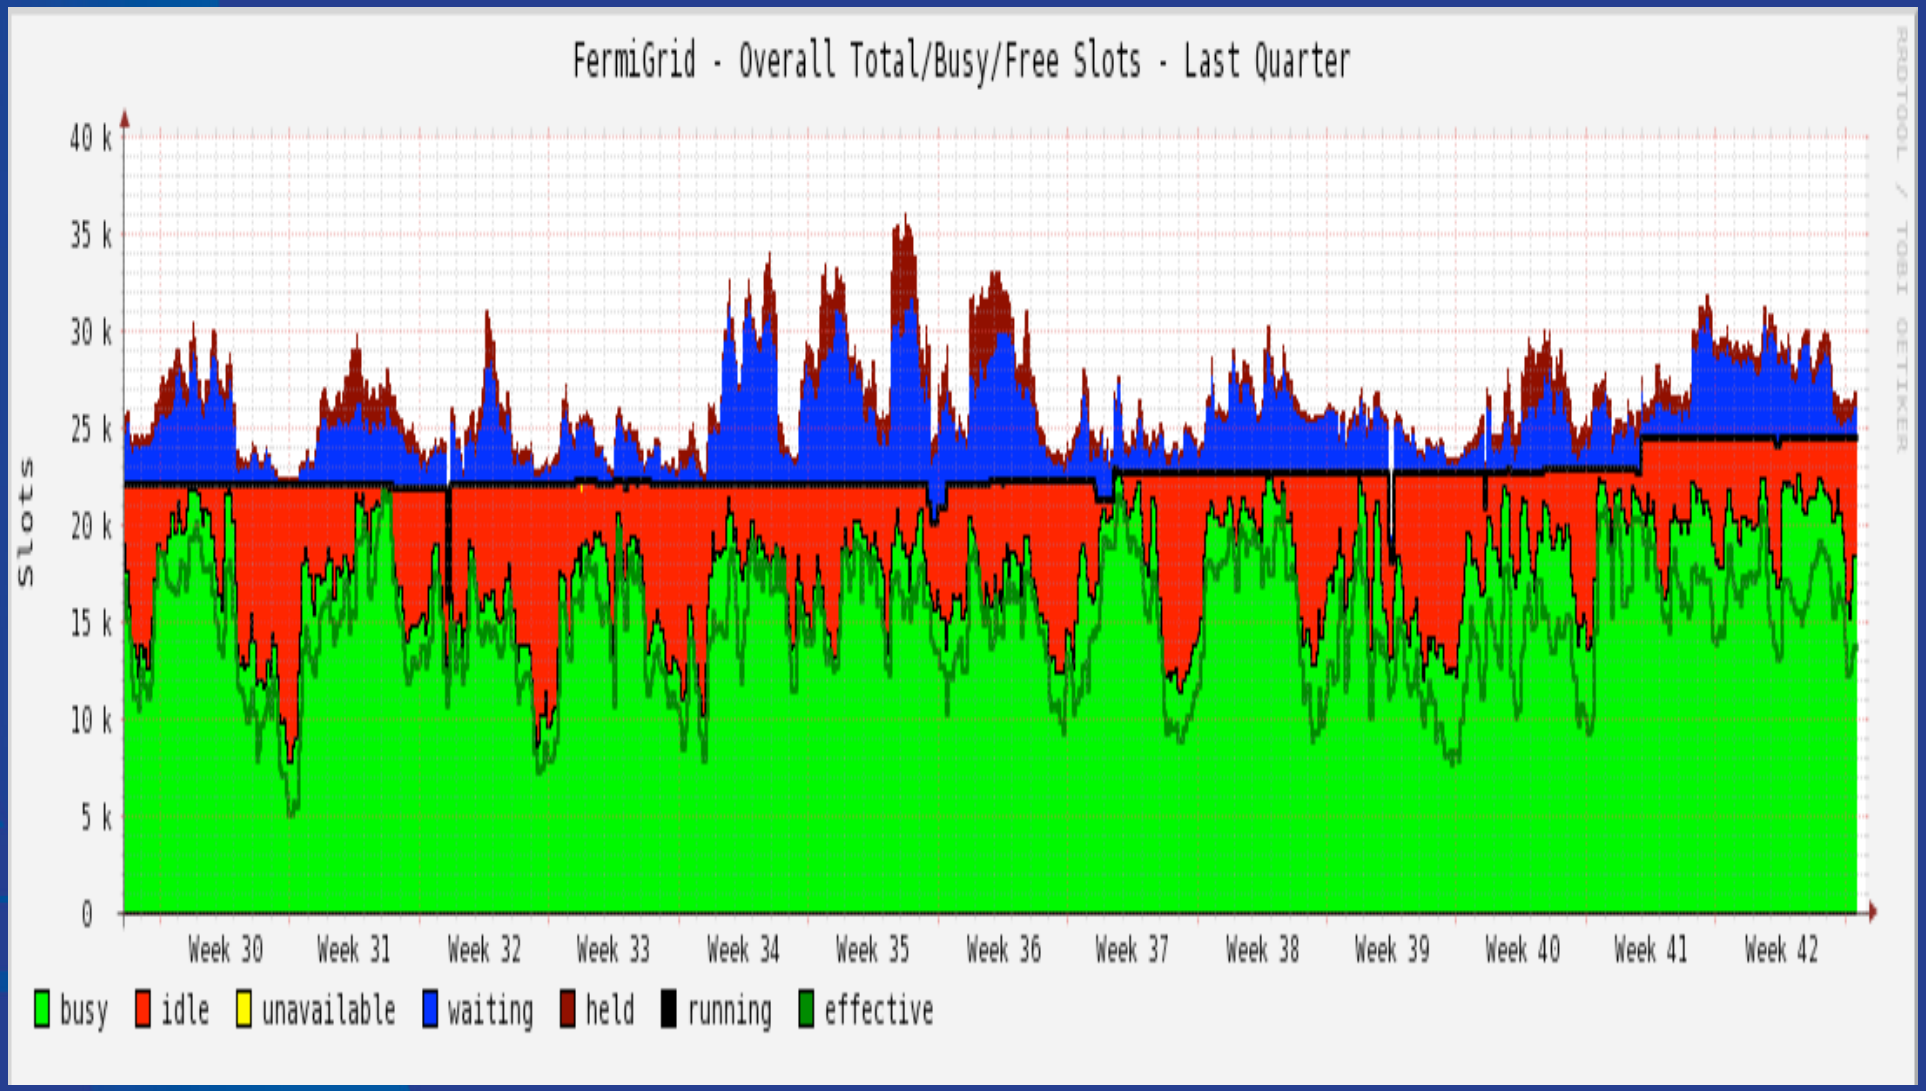

FermiGrid Service Availability (measured over the past year)

Service	Raw Availability	HA Configuration	Measured HA Availability	Minutes of Downtime
VOMS – VO Management Service	99.977	Active-Active	99.988	60
GUMS – Grid User Mapping Service	99.989	Active-Active	100.000	0
SAZ – Site AuthoriZation Service	99.989	Active-Active	100.000	0
Squid – Web Cache	99.793	Active-Active	100.000	0
MyProxy – Grid Proxy Service	99.499	Active-Standby	99.954	240
ReSS – Resource Selection Service	99.966	Active-Active	99.977	120
Gratia – Fermilab and OSG Accounting	97.914	Active-Standby	99.945	300
MySQL Database	98.054	Active-Active	99.988	60

FermiCloud

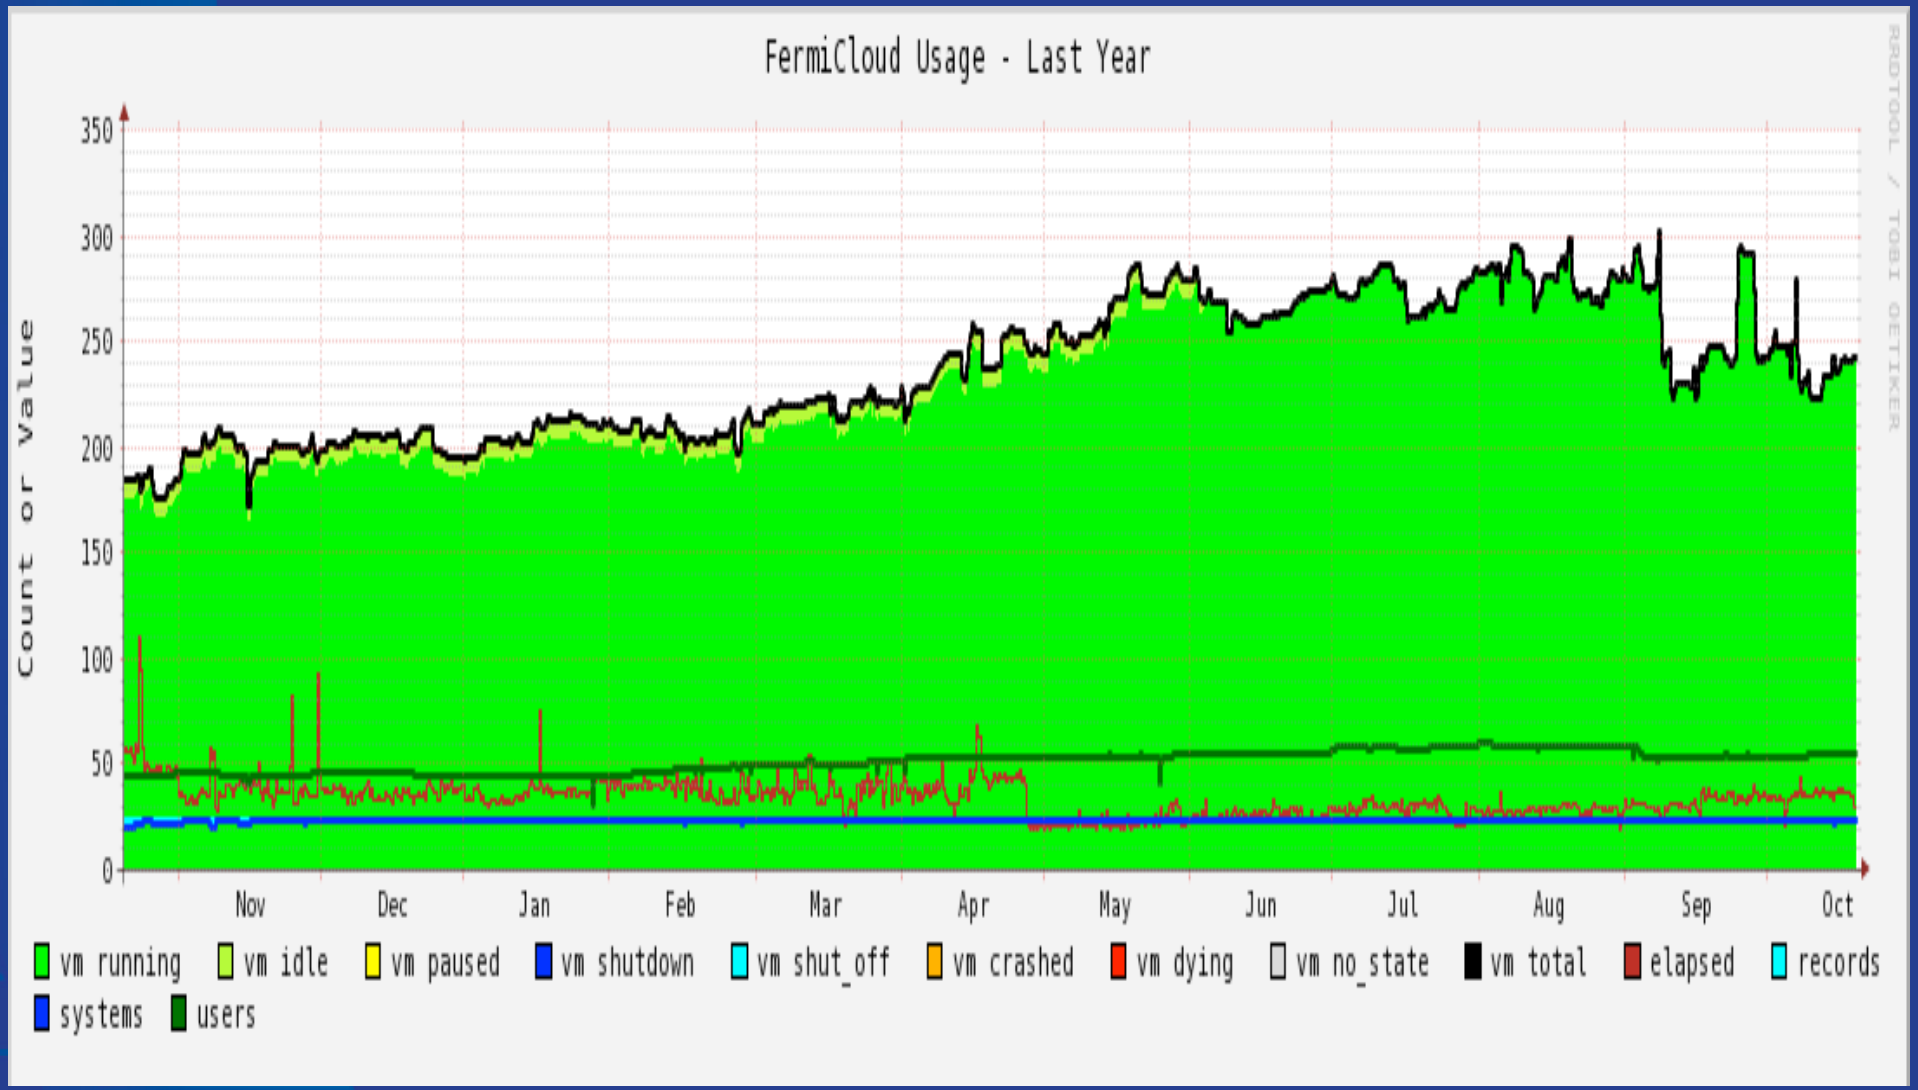

For further details see Gerard's talk later in the week...

FIFE
(FabrIc for Frontier Experiments)

A collaborative effort, led by the Computing Sector, to take the lessons learned from supporting Run II and LHC experiments and appropriately incorporate them into the software frameworks being used by the Frontier Experiments:

- FIFE Architecture Report:
  - http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5180
- IFDH – Transparently "moving" data access from BlueArc to dCache,
- Jobsub – A client-server utility to submit jobs on behalf of the various Frontier experiments in a GlideinWMS framework,
- Significant investment in use of CVMFS to support publication of the "standard" analysis packages both onsite and offsite.

Energy Frontier - CDF & D0

CDF and D0 continue to perform analysis of their datasets,

The Run II Data Preservation efforts continue to gather steam...

Work is underway to replicate the CDF dataset from Fermilab to INFN.

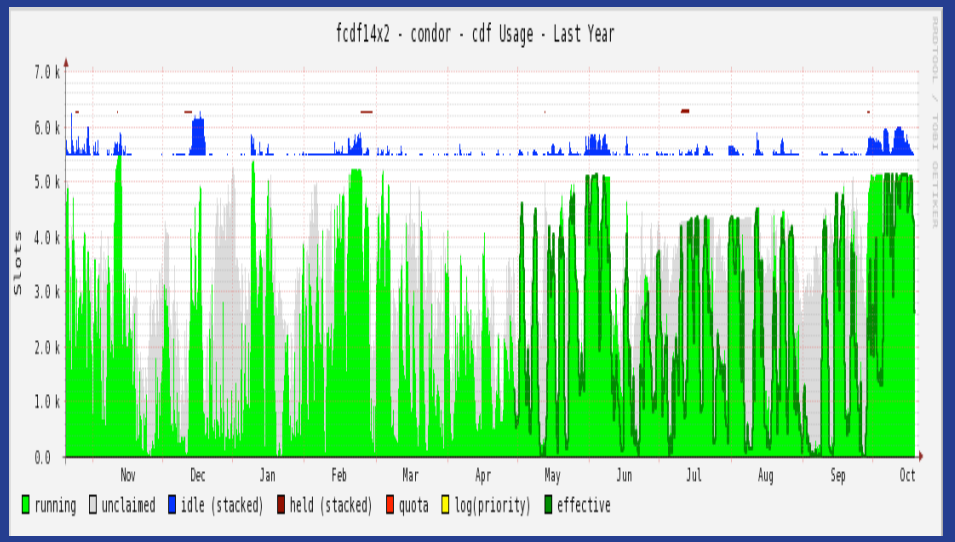

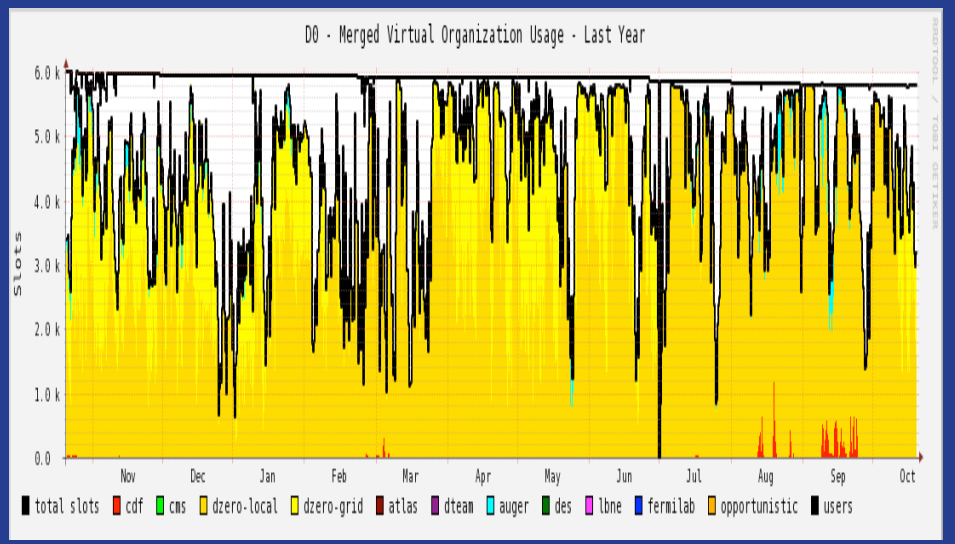

Energy Frontier - CMS

The CMS T1 at Fermilab continues to deliver computing to the global CMS analysis effort:

- Availability is well above the pledge requirements,
- Fermilab is regularly at or near the top of the availability rank plot.

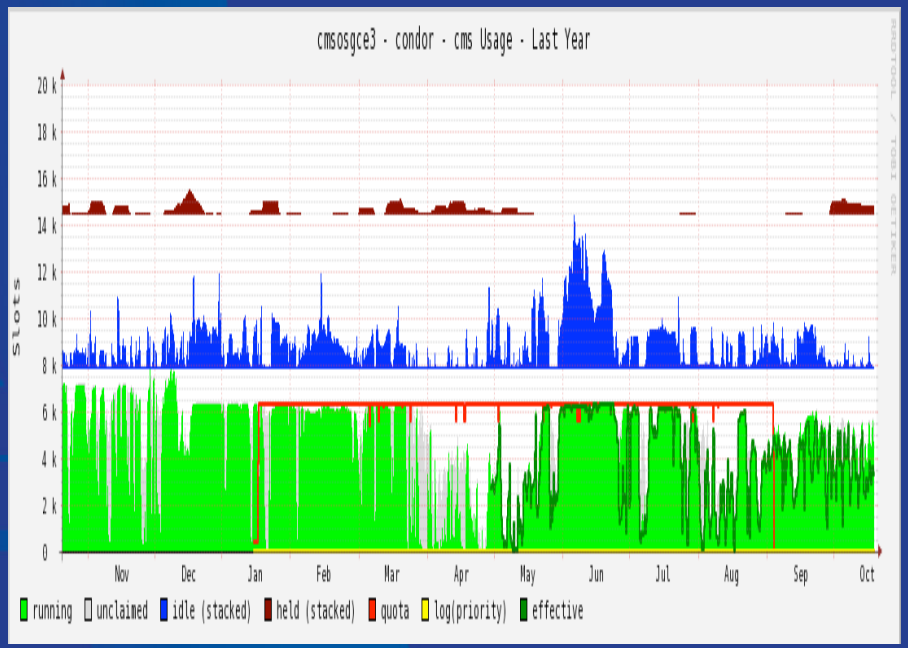

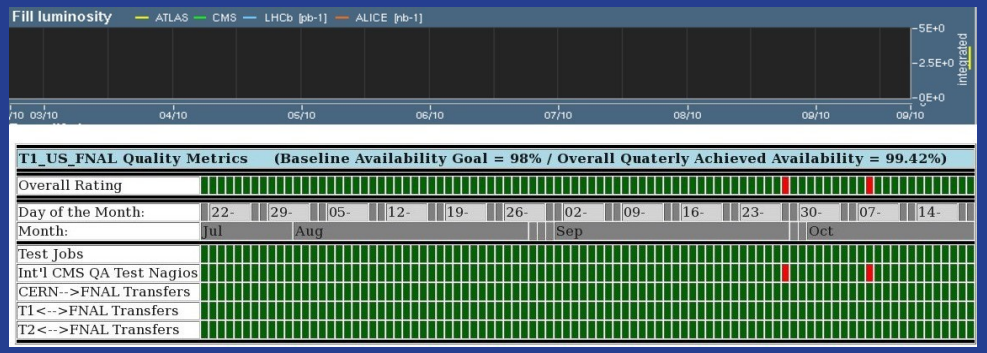

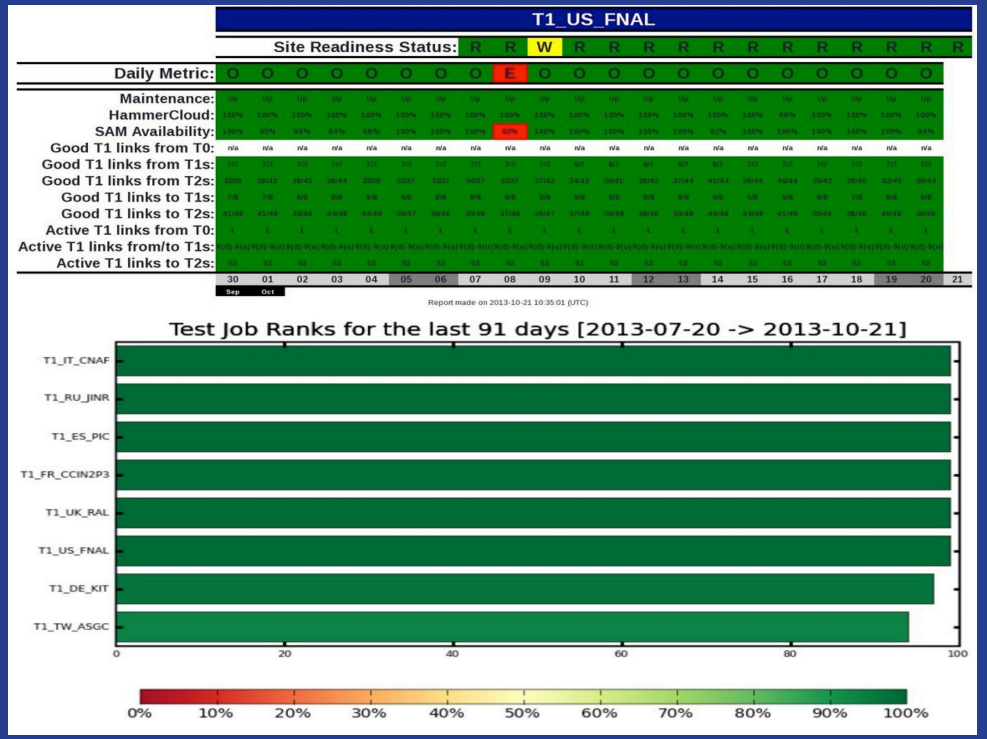

Intensity Frontier
NOvA Far Detector Assembly Progress

NOvA Webcam:

http://www.fnal.gov/pub/webcams/nova_webcam/

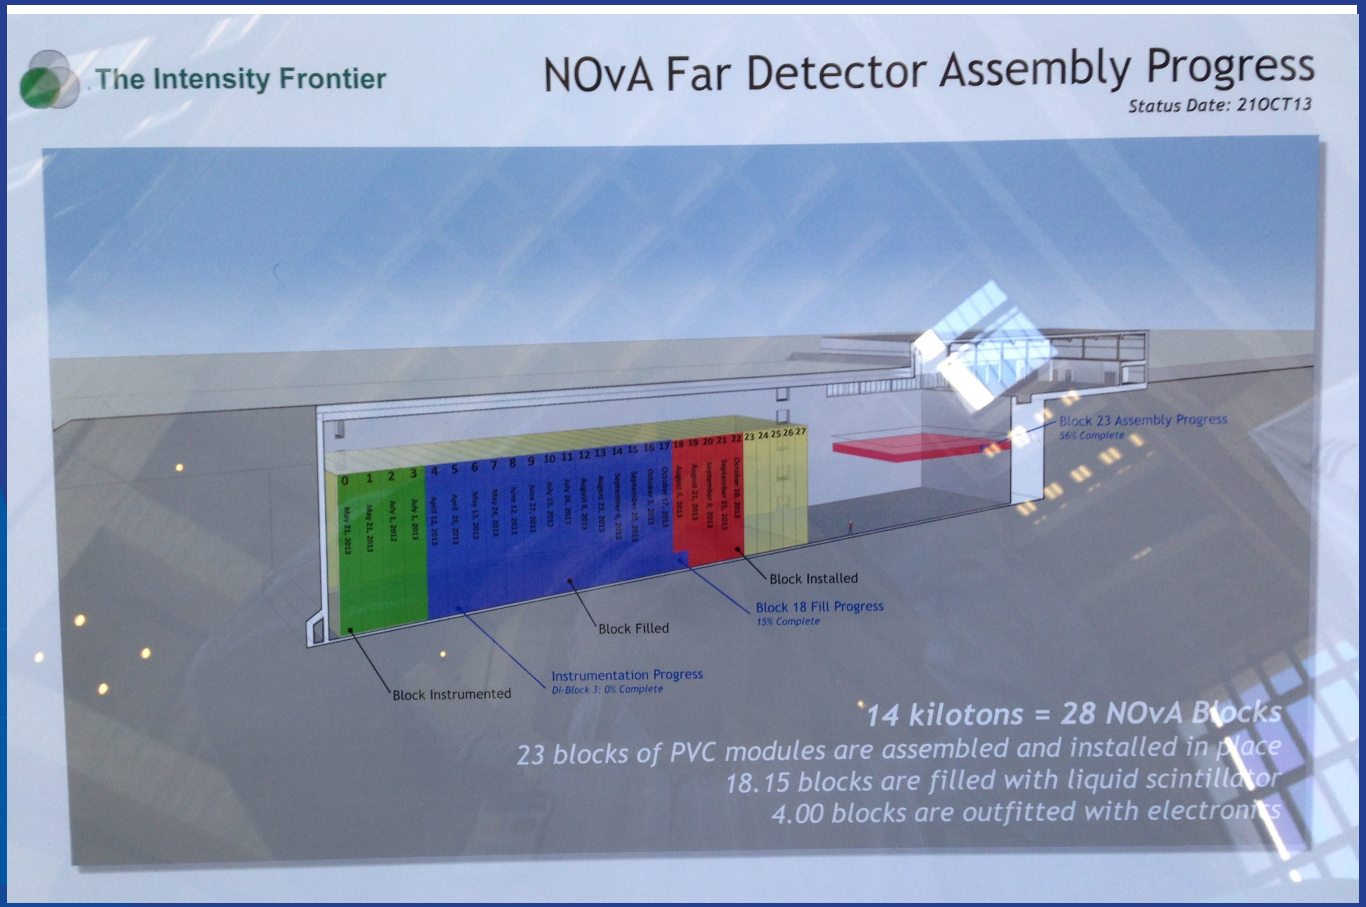

Cosmic Frontier - DES

Dark Energy Survey:

- The Dark Energy Survey officially began on August 31, 2013,
- Using the powerful Dark Energy Camera, scientists will map a portion of the sky in unprecedented detail, seeking answers to the mystery of dark energy.
- http://www.fnal.gov/pub/presspass/press_releases/2013/DES-20130903.html

More information:

- https://www.darkenergysurvey.org/

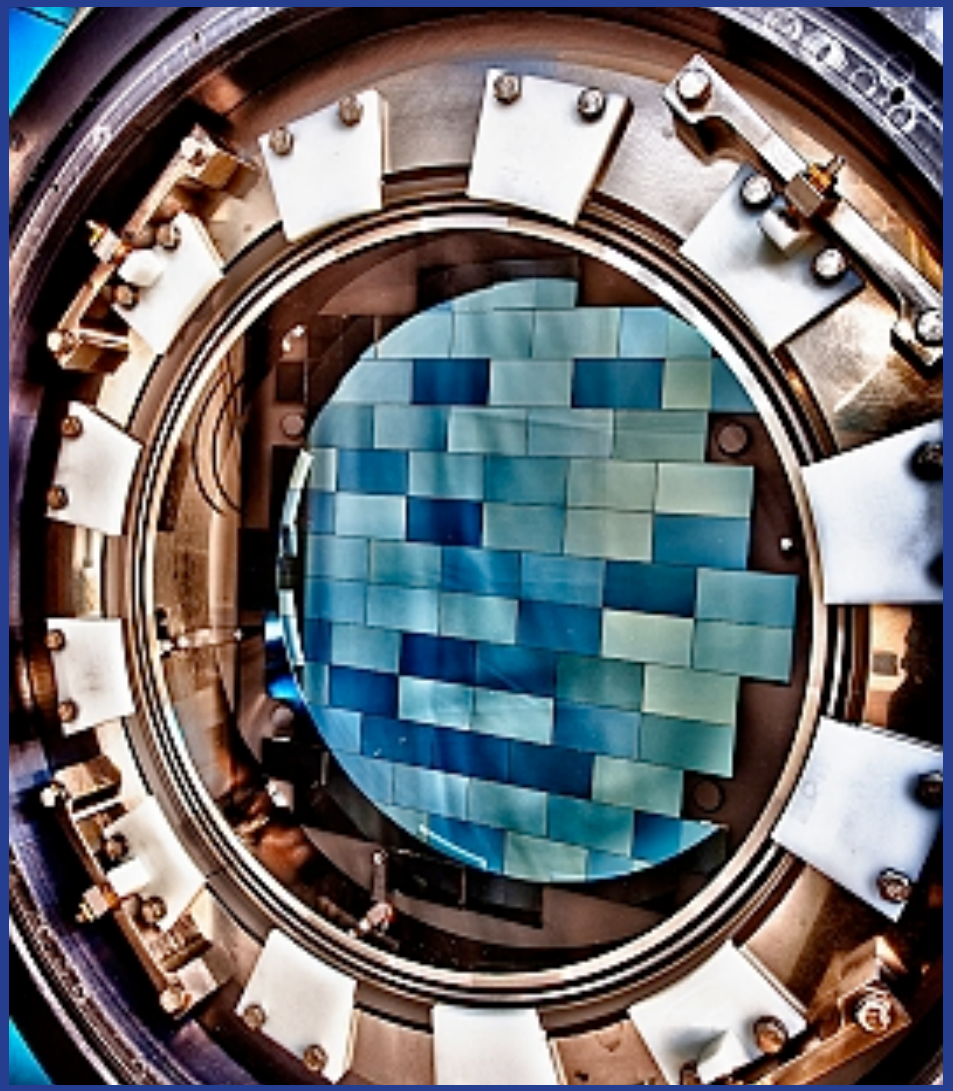

Illinois Accelerator Research Center

Other talks from Fermilab personnel this week…

Presenter	Presentation Title
Kevin Hill	Evolution of the OSG Authentication Model
Tyler Parsons & Edward Simmonds	Puppet at Fermilab- Managing a Large Heterogeneous Environment
Tim Skirvin	Puppet at USCMS-T1 and FermiLab
Lisa Giacchetti	USCMS T1 and LPC Data Storage Challenges and Solutions
Alexander Moibenko	Solving Small Files Problem in Enstore
Bonnie King, Pat Riehecky & Connie Sieh	Scientific Linux current status update
Gerard Bernabeu	FermiCloud update - enabling scientific workflows in the cloud

Summary

It is an exciting time to be supporting science at Fermilab in all of the three frontiers of science:

- Energy Frontier
- Intensity Frontier
- Cosmic Frontier

And there is no danger of running out of work supporting and extending the computer systems and services that provide the foundations of supporting science.

Thank You!

Any Questions?

Extra Slides

Office of the CIO
Vicky White

New Deputy CIO - Jin Chang (formerly CIO of NYC Department of the Environment).

Eight "Groups":

- Governance – Bill Boroski
- Administrative Support – Bernadette Tabor
- Computer Security – Joe Klemencic
- Enterprise Architecture – Scott Nolan
- Financial Management – Valena Sibley
- Project Management - Bill Boroski
- Service Management – Tammy Whited
- Communications – Marcia Teckenbrock

Core Computing Division
Jon Bakken

Six Departments:

- Cyber Security Services - Joe Klemencic
- Service Operations Support – Eileen Berman
- Information Systems - Santo Campione
- Network and Communication Services – Ramon Pasetes
- Enterprise Services Operations - Michael Rosier
- Facility Operations – Adam Walters

Scientific Computing Division
Rob Roser

Two groups:

- Scientific Technical Architecture - Ruth Pordes
- Entrepreneurial Ventures - Rob Roser

Four quadrants (16 Departments):

- Systems for Scientific Applications - Panagiotis Spentzouris
  - Scientific Software Infrastructure, Electronic Systems Engineering, Scientific Computing Simulation, Real-Time Software Infrastructure
- Scientific Programs - Lothar Bauerdick
  - CMS, Theory, Emeritus Scientists, Intensity Frontier, Experimental Astrophysics
- Scientific Computing Services - Margaret Votava
  - Grid and Cloud Services, Scientific Data Processing
- Scientific Computing Facilities - Stuart Fuess
  - High Performance Parallel Computing Facilities, CMS Computing Facilities, Data Movement And Storage, Fermilab Experiments Facilities

New SCD Virtualization and Cloud Strategy Goals[1]

Strategic Goals:

1. The Laboratory's scientific mission and Enterprise Architecture are overarching principles guiding the choice of solutions.
2. Operate cost-effective, usable and expandable Scientific Computing services and resources to meet our users needs for the next five to ten years.
3. Position our scientific computing solutions to take advantage of current and future projections and opportunities in an environment of rapidly changing technologies and cost opportunities.
4. Ensure that our scientific users can access and use with agility the appropriate local, remote, opportunistic, commercial, High Performance Computing and High Throughput Computing resources to meet their needs with a minimum of overheads.
5. Provide an environment for scientific computing applications that can integrate and adapt to new technologies, methods, and the evolution in resources as needed easy to use and a good fit for the people who use it.
6. including DOE and NSF leadership class machines.

New SCD Virtualization and Cloud Strategy [2]

Strategies:

1. Deploy common solutions wherever it makes sense.
2. Provision resources and services as virtualized environments unless it does not make sense.
3. Manage and support dynamic provisioning of resources and services (based on agreed upon policies), including support for "opportunistic needs", whenever it makes sense.
4. Ensure smooth transitions to new technologies and, where beneficial, collaborate in the development of the underlying technologies to meet our strategic goals.
5. Support more than one set of technologies for scientific users, including a legacy system, production solution, and forward looking but perhaps less robust solution.
6. Let cost effectiveness be a strong guide to choice of solution and technologies. SCD will publish and evolve the process for determination of "makes sense".

New SCD Virtualization and Cloud Strategy [3]

Assumptions:

1. User requirements and technology performance will dictate a mix of resource performance and types – including support for high performance I/O, execution of some jobs on "bare metal", high-availability, and virtualized services and resources. The balance will be determined by a combination of drivers from the users, stakeholders and the laboratory.
2. If heterogeneous solutions are implemented the deployment of common interfaces across them will be a priority.
3. ITIL based processes and service management will be implemented as part of any/all solutions.

General Physics Compute Facility (GPCF)

GPCF provides statically deployed virtual machines:

- GPCF is a critical component of the Intensity and Cosmic Frontier experiments,
- As well as others who need similar services,
- It continues to operate well.

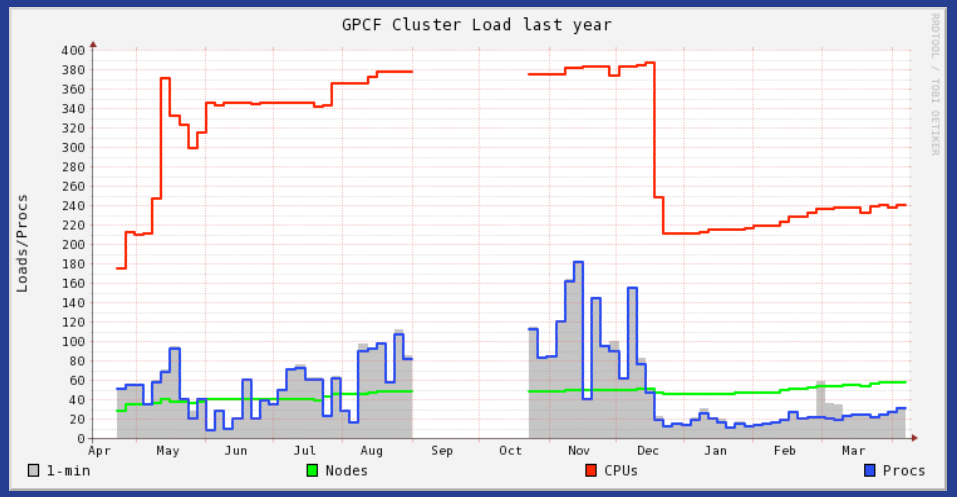

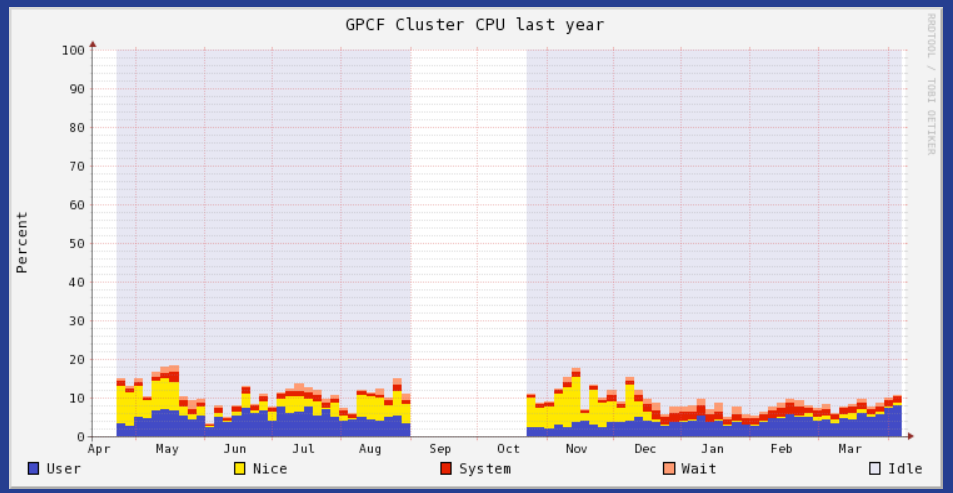

VMware Virtual Services

Primary focus is the support of Fermilab core computing division services,